


<h1>TCSS 422: OPERATING SYSTEMS</h1>	
<p>Introduction to Concurrency, Locks, Lock-based Data Structures, Condition Variables</p> <p>Wes J. Lloyd School of Engineering and Technology University of Washington - Tacoma</p> <p>February 6, 2019</p>	
<p>TCSS422: Operating Systems [Winter 2019] School of Engineering and Technology, University of Washington - Tacoma</p>	

<h2>OBJECTIVES</h2>		
<ul style="list-style-type: none">■ C Tutorial■ Assignment 1■ Midterm 2/13■ Feedback 1/30 ■ <u>Parallel programming with P-threads cont'd</u>■ Chapter 27 – Linux Thread API■ Chapter 28 – Intro to locks■ Chapter 29 – Lock-Based Data Structures■ Chapter 30 – Condition Variables		
February 6, 2019	TCSS422: Operating Systems [Winter 2019] School of Engineering and Technology, University of Washington - Tacoma	L8.2

FEEDBACK FROM 1/30

- **Ticket Distribution – Proportional Share Schedulers:**
- ***Can nice values be used to determine ticket distribution?***
- **Linux nice values:**
- Range from: -20 (highest priority) to 19 (lowest priority)
- Can't directly use the nice value to assign tickets
- Job with -20 tickets !!!
- CAN use nice value to identify jobs with the same priority
 - Assign tickets proportionally for jobs with same priority
- Need to determine how many tickets each priority level should receive to share amongst its jobs
 - Will vary based on # of jobs at the priority level

February 6, 2019

TCSS422: Operating Systems [Winter 2019]
School of Engineering and Technology, University of Washington - Tacoma

L8.3

FEEDBACK - 2

- **Didn't quite understand parallel programming and locks**
- **What is parallel programming?**
 - Parallel programming – use of multiple threads to execute over the same program code at the same time sharing memory
- **What data is shared by threads?**
 - Heap segment, data segment (global variables), code segment
- **What do locks do?**
 - Locks BLOCK multiple threads from executing ***critical sections*** of code at the same time, making execution ***atomic*** within these sections
- **What is a blocking API call?**
 - A kernel function that “hibernates” the user thread to wait for a resource to become available. The users thread goes from RUNNING→BLOCKED. When the resource is available, the OS generates an interrupt, and the user thread is awoken to process the interrupt.
- **Is pthread_mutex_lock() a blocking API call?**
 - YES – Note that if multiple threads are sleeping for the lock, only one gets woken up – this is chosen by the kernel – fairness can be an issue

February 6, 2019

TCSS422: Operating Systems [Winter 2019]
School of Engineering and Technology, University of Washington - Tacoma

L8.4

FEEDBACK - 3

- Can you run an entire program with atomic execution?
- Good question!, SURE, there is no reason this wouldn't be allowed, --BUT-- this scenario may have little value as essentially the program would become sequential, and operate as if "single threaded"
-- nothing can happen in parallel!
- Does C or any other high level programming language automatically create multiple threads for a process?
- SURE, high level languages may include functions or classes that automatically create worker threads to complete tasks in parallel
- Example: Java `Array.parallelSort()` -- added in Java 8

February 6, 2019

TCS422: Operating Systems [Winter 2019]
School of Engineering and Technology, University of Washington - Tacoma

L8.5

JAVA ARRAY PARALLEL SORT EXAMPLE

```
import java.util.Arrays;
public class Example
{
    public static void main(String[] args)
    {
        int numbers[] = {22, 89, 1, 32, 19, 5};
        //Parallel Sort method for sorting int array
        ➡ Arrays.parallelSort(numbers);
        //convert array to stream and display w/
        forEach
        Arrays.stream(numbers).forEach(n->
        System.out.print(n+" "));
    }
}
```

February 6, 2019

TCS422: Operating Systems [Winter 2019]
School of Engineering and Technology, University of Washington - Tacoma

L8.6

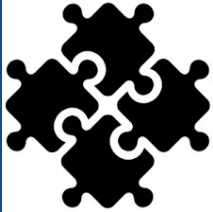
FEEDBACK - 4

- What chapters / subjects will the midterm cover?
- Midterm Wednesday February 13th
- Inclusive of content covered in class through February 11th

February 6, 2019	TCSS422: Operating Systems [Winter 2019] School of Engineering and Technology, University of Washington - Tacoma	L8.7
------------------	---	------

CHAPTER 27 - LINUX THREAD API

February 6, 2019	TCSS422: Operating Systems [Winter 2019] School of Engineering and Technology, University of Washington - Tacoma	L8.8
------------------	---	------



THREAD CREATION

■ pthread_create

```
#include <pthread.h>

int
pthread_create(      pthread_t*      thread,
                    const pthread_attr_t* attr,
                    void*             (*start_routine) (void*),
                    void*             arg);
```

- thread: thread struct
- attr: stack size, scheduling priority... (*optional*)
- start_routine: function pointer to thread routine
- arg: argument to pass to thread routine (*optional*)

February 6, 2019

TCSS422: Operating Systems [Winter 2019]
School of Engineering and Technology, University of Washington - Tacoma

L8.9

PTHREAD_CREATE – PASS ANY DATA

```
#include <pthread.h>

typedef struct __myarg_t {
    int a;
    int b;
} myarg_t;

void *mythread(void *arg) {
    myarg_t *m = (myarg_t *) arg;
    printf("%d %d\n", m->a, m->b);
    return NULL;
}

int main(int argc, char *argv[]) {
    pthread_t p;
    int rc;

    myarg_t args;
    args.a = 10;
    args.b = 20;
    rc = pthread_create(&p, NULL, mythread, &args);
    ...
}
```

February 6, 2019

TCSS422: Operating Systems [Winter 2019]
School of Engineering and Technology, University of Washington - Tacoma

L8.10

PASSING A SINGLE VALUE

Using this approach on your Ubuntu VM,
How large (in bytes) can the primitive data type be?

How large (in bytes) can the primitive data type
be on a 32-bit operating system?

```
9   int rc, m;  
10  pthread_create(&p, NULL, mythread, (void *) 100);  
11  pthread_join(p, (void **) &m);  
12  printf("returned %d\n", m);  
13  return 0;  
14 }
```

February 6, 2019

TCSS422: Operating Systems [Winter 2019]
School of Engineering and Technology, University of Washington - Tacoma

L8.11

WAITING FOR THREADS TO FINISH

```
int pthread_join(pthread_t thread, void **value_ptr);
```

- thread: which thread?
- value_ptr: pointer to return value
type is dynamic / agnostic
- Returned values **must** be on the heap
- Thread stacks destroyed upon thread termination (join)
- Pointers to thread stack memory addresses are invalid
 - May appear as gibberish or lead to crash (seg fault)
- Not all threads join – *What would be Examples ??*

February 6, 2019

TCSS422: Operating Systems [Winter 2019]
School of Engineering and Technology, University of Washington - Tacoma

L8.12

```

struct myarg {
    int a;
    int b;
};

void *worker(void *arg)
{
    struct myarg *input = (struct myarg *) arg;
    printf("a=%d b=%d\n", input->a, input->b);
    struct myarg output;
    output.a = 1;
    output.b = 2;
    return (void *) &output;
}

int main (int argc, char * argv[])
{
    pthread_t p1;
    struct myarg args;
    struct myarg *ret_args;
    args.a = 10;
    args.b = 20;
    pthread_create(&p1, NULL, worker, &args);
    pthread_join(p1, (void **)&ret_args);
    printf("returned %d %d\n", ret_args->a, ret_args->b);
    return 0;
}

```

What will this code do?

Data on thread stack

```

$ ./pthread_struct
a=10 b=20
Segmentation fault (core dumped)

```

How can this code be fixed?

February 6, 2019
TCSS422: Operating Systems [Winter 2019]
School of Engineering and Technology, University of Washington - Tacoma
L8.13

```

struct myarg {
    int a;
    int b;
};

void *worker(void *arg)
{
    struct myarg *input = (struct myarg *) arg;
    printf("a=%d b=%d\n", input->a, input->b);
    input->a = 1;
    input->b = 2;
    return (void *) &input;
}

int main (int argc, char * argv[])
{
    pthread_t p1;
    struct myarg args;
    struct myarg *ret_args;
    args.a = 10;
    args.b = 20;
    pthread_create(&p1, NULL, worker, &args);
    pthread_join(p1, (void **)&ret_args);
    printf("returned %d %d\n", ret_args->a, ret_args->b);
    return 0;
}

```

How about this code?

```

$ ./pthread_struct
a=10 b=20
returned 1 2

```

February 6, 2019
TCSS422: Operating Systems [Winter 2019]
School of Engineering and Technology, University of Washington - Tacoma
L8.14

ADDING CASTS

- Casting
- Suppresses compiler warnings when passing “typed” data where (void) or (void *) is called for
- Example: uncasted capture in pthread_join
pthread_int.c: In function 'main':
pthread_int.c:34:20: warning: passing argument 2 of 'pthread_join' from incompatible pointer type [-Wincompatible-pointer-types]
pthread_join(p1, &p1val);
- Example: uncasted return
In file included from pthread_int.c:3:0:
/usr/include/pthread.h:250:12: note: expected 'void **' but argument is of type 'int **'
extern int pthread_join (pthread_t __th, void **__thread_return);

February 6, 2019

TCSS422: Operating Systems [Winter 2019]
School of Engineering and Technology, University of Washington - Tacoma

L8.15

ADDING CASTS - 2

- pthread_join
int * p1val;
int * p2val;
pthread_join(p1, (void *)&p1val);
pthread_join(p2, (void *)&p2val);
- return from thread function
int * counterval = malloc(sizeof(int));
*counterval = counter;
return (void *) counterval;

February 6, 2019

TCSS422: Operating Systems [Winter 2019]
School of Engineering and Technology, University of Washington - Tacoma

L8.16

LOCKS

- `pthread_mutex_t` data type
- `/usr/include/bits/pthread_types.h`

```
// Global Address Space
static volatile int counter = 0;
pthread_mutex_t lock;

void *worker(void *arg)
{
    int i;
    for (i=0;i<10000000;i++) {
        int rc = pthread_mutex_lock(&lock);
        assert(rc==0);
        counter = counter + 1;
        pthread_mutex_unlock(&lock);
    }
    return NULL;
}
```

February 6, 2019

TCSS422: Operating Systems [Winter 2019]
School of Engineering and Technology, University of Washington - Tacoma

L8.17

LOCKS - 2

- Ensure critical sections are executed atomically-as a *unit*
 - Provides implementation of “*Mutual Exclusion*”

- API

```
int pthread_mutex_lock(pthread_mutex_t *mutex);
int pthread_mutex_unlock(pthread_mutex_t *mutex);
```

- Example w/o initialization & error checking



```
pthread_mutex_t lock;
pthread_mutex_lock(&lock);
x = x + 1; // or whatever your critical section is
pthread_mutex_unlock(&lock);
```

- Blocks forever until lock can be obtained
- Enters critical section once lock is obtained
- Releases lock

February 6, 2019

TCSS422: Operating Systems [Winter 2019]
School of Engineering and Technology, University of Washington - Tacoma

L8.18

LOCK INITIALIZATION

- Assigning the constant

```
pthread_mutex_t lock = PTHREAD_MUTEX_INITIALIZER;
```

- API call:

```
int rc = pthread_mutex_init(&lock, NULL);  
assert(rc == 0); // always check success!
```

- Initializes mutex with attributes specified by 2nd argument
- If NULL, then default attributes are used
- Upon initialization, the mutex is initialized and unlocked

February 6, 2019

TCSS422: Operating Systems [Winter 2019]
School of Engineering and Technology, University of Washington - Tacoma

L8.19

LOCKS - 3

- Error checking wrapper

```
// Use this to keep your code clean but check for failures  
// Only use if exiting program is OK upon failure  
void Pthread_mutex_lock(pthread_mutex_t *mutex) {  
    int rc = pthread_mutex_lock(mutex);  
    assert(rc == 0);  
}
```

- What if lock can't be obtained?

```
int pthread_mutex_trylock(pthread_mutex_t *mutex);  
int pthread_mutex_timelock(pthread_mutex_t *mutex,  
                           struct timespec *abs_timeout);
```

- trylock – returns immediately (fails) if lock is unavailable
- timelock – tries to obtain a lock for a specified duration

February 6, 2019

TCSS422: Operating Systems [Winter 2019]
School of Engineering and Technology, University of Washington - Tacoma

L8.20

CONDITIONS AND SIGNALS

- Condition variables support “signaling” between threads

```
int pthread_cond_wait(pthread_cond_t *cond,  
                      pthread_mutex_t *mutex);  
int pthread_cond_signal(pthread_cond_t *cond);
```



- pthread_cond_t datatype

- pthread_cond_wait()
 - Puts thread to “sleep” (waits) (THREAD is BLOCKED)
 - Threads added to FIFO queue, lock is released
 - Waits (*listens*) for a “signal” (NON-BUSY WAITING, no polling)
 - When signal occurs, interrupt fires, wakes up first thread, (THREAD is RUNNING), lock is provided to thread

February 6, 2019

TCSS422: Operating Systems [Winter 2019]
School of Engineering and Technology, University of Washington - Tacoma

L8.21

CONDITIONS AND SIGNALS - 2

```
int pthread_cond_signal(pthread_cond_t * cond);  
int pthread_cond_broadcast(pthread_cond_t * cond);
```

- pthread_cond_signal()
 - Called to send a “signal” to wake-up first thread in FIFO “wait” queue
 - The goal is to unblock a thread to respond to the signal
- pthread_cond_broadcast()
 - Unblocks *all* threads in FIFO “wait” queue, currently blocked on the specified condition variable
 - Broadcast is used when all threads should wake-up for the signal
- Which thread is unblocked first?
 - Determined by OS scheduler (based on priority)
 - Thread(s) awoken based on placement order in FIFO wait queue
 - When awoken threads acquire lock as in pthread_mutex_lock()

February 6, 2019

TCSS422: Operating Systems [Winter 2019]
School of Engineering and Technology, University of Washington - Tacoma

L8.22

CONDITIONS AND SIGNALS - 3

■ Wait example:

```
pthread_mutex_t lock = PTHREAD_MUTEX_INITIALIZER;
pthread_cond_t cond = PTHREAD_COND_INITIALIZER;

pthread_mutex_lock(&lock);
while (initialized == 0)
    pthread_cond_wait(&cond, &lock);
// Perform work that requires lock
a = a + b;
pthread_mutex_unlock(&lock);
```

- wait puts thread to sleep, releases lock
- when awoken, lock reacquired (but then released by this code)
- When initialized, another thread signals

```
pthread_mutex_lock(&lock);
initialized = 1;
pthread_cond_signal(&cond);
pthread_mutex_unlock(&lock);
```

State variable set,
Enables other thread(s)
to proceed above.

February 6, 2019

TCSS422: Operating Systems [Winter 2019]
School of Engineering and Technology, University of Washington - Tacoma

L8.23

CONDITION AND SIGNALS - 4

```
pthread_mutex_t lock = PTHREAD_MUTEX_INITIALIZER;
pthread_cond_t cond = PTHREAD_COND_INITIALIZER;

pthread_mutex_lock(&lock);
while (initialized == 0)
    pthread_cond_wait(&cond, &lock);
// Perform work that requires lock
a = a + b;
pthread_mutex_unlock(&lock);
```

- Why do we wait inside a while loop?
- The while ensures upon awakening the condition is rechecked
 - A signal is raised, but the pre-conditions required to proceed may have not been met. ****MUST CHECK STATE VARIABLE****
 - Without checking the state variable the thread may proceed to execute when it should not. (e.g. too early)

February 6, 2019

TCSS422: Operating Systems [Winter 2019]
School of Engineering and Technology, University of Washington - Tacoma

L8.24

PTHREADS LIBRARY

- **Compilation**
 - `gcc -pthread pthread.c -o pthread`
 - Requires explicitly linking the library with compiler flag
 - Use makefile to provide compiler arguments
- **List of pthread manpages**
 - `man -k pthread`

February 6, 2019

TCSS422: Operating Systems [Winter 2019]
School of Engineering and Technology, University of Washington - Tacoma

L8.25

SAMPLE MAKEFILE

```
CC=gcc
CFLAGS=-pthread -I. -Wall

binaries=pthread pthread_int pthread_lock_cond pthread_struct
all: $(binaries)

pthread_mult: pthread.c pthread_int.c
    $(CC) $(CFLAGS) $^ -o $@

clean:
    $(RM) -f $(binaries) *.o
```

- **Example builds multiple single file programs**
 - All target
- **pthread_mult**
 - Example if multiple source files should produce a single executable
- **clean target**

February 6, 2019

TCSS422: Operating Systems [Winter 2019]
School of Engineering and Technology, University of Washington - Tacoma

L8.26

CHAPTER 28 – LOCKS

February 6, 2019

TCSS422: Operating Systems [Winter 2019]
School of Engineering and Technology, University of Washington - Tacoma

L8.27



LOCKS



- Ensure critical section(s) are executed atomically-as a *unit*
 - Only one thread is allowed to execute a critical section at any given time
 - Ensures the code snippets are “mutually exclusive”

- Protect a global counter:

```
balance = balance + 1;
```

- A “critical section”:

```
1  lock_t mutex; // some globally-allocated lock 'mutex'
2  ...
3  lock(&mutex);
4  balance = balance + 1;
5  unlock(&mutex);
```

February 6, 2019

TCSS422: Operating Systems [Winter 2019]
School of Engineering and Technology, University of Washington - Tacoma

L8.28

LOCKS - 2

- Lock variables are called “MUTEX”
 - Short for mutual exclusion (that’s what they guarantee)
- Lock variables store the state of the lock
- States
 - **Locked** (acquired or held)
 - **Unlocked** (available or free)
- Only 1 thread can hold a lock

February 6, 2019

TCSS422: Operating Systems [Winter 2019]
School of Engineering and Technology, University of Washington - Tacoma

L8.29

LOCKS - 3

- `pthread_mutex_lock(&lock)`
 - Try to acquire lock
 - If lock is free, calling thread will acquire the lock
 - Thread with lock enters critical section
 - Thread “owns” the lock
- No other thread can acquire the lock before the owner releases it.

February 6, 2019

TCSS422: Operating Systems [Winter 2019]
School of Engineering and Technology, University of Washington - Tacoma

L8.30

LOCKS - 4

- Program can have many mutex (lock) variables to “serialize” many critical sections
- Locks are also used to protect data structures
 - Prevent multiple threads from changing the same data simultaneously
 - Programmer can make sections of code “granular”
 - Fine grained – means just one grain of sand at a time through an hour glass
 - Similar to relational database transactions
 - DB transactions prevent multiple users from modifying a table, row, field

February 6, 2019


TCSS422: Operating Systems [Winter 2019]
School of Engineering and Technology, University of Washington - Tacoma

L8.31

FINE GRAINED?

- Is this code a good example of “fine grained parallelism”?

```
pthread_mutex_lock(&lock);  
a = b++;  
b = a * c;  
*d = a + b + c;  
FILE * fp = fopen ("file.txt", "r");  
fscanf(fp, "%s %s %s %d", str1, str2, str3, &e);  
ListNode *node = mylist->head;  
Int i=0  
while (node) {  
    node->title = str1;  
    node->subheading = str2;  
    node->desc = str3;  
    node->end = *e;  
    node = node->next;  
    i++  
}  
e = e - i;  
pthread_mutex_unlock(&lock);
```



February 6, 2019

TCSS422: Operating Systems [Winter 2019]
School of Engineering and Technology, University of Washington - Tacoma

L8.32

FINE GRAINED PARALLELISM

```
pthread_mutex_lock(&lock_a);
pthread_mutex_lock(&lock_b);
a = b++;
pthread_mutex_unlock(&lock_b);
pthread_mutex_unlock(&lock_a);

pthread_mutex_lock(&lock_b);
b = a * c;
pthread_mutex_unlock(&lock_b);

pthread_mutex_lock(&lock_d);
*d = a + b + c;
pthread_mutex_unlock(&lock_d);

FILE * fp = fopen ("file.txt", "r");
pthread_mutex_lock(&lock_e);
fscanf(fp, "%s %s %s %d", str1, str2, str3, &e);
pthread_mutex_unlock(&lock_e);

ListNode *node = mylist->head;
int i=0 . . .
```



February 6, 2019

TCCS422: Operating Systems [Winter 2019]
School of Engineering and Technology, University of Washington - Tacoma

L8.33

EVALUATING LOCK IMPLEMENTATIONS

■ Correctness

- Does the lock work?
- Are critical sections mutually exclusive?
(atomic-as a unit?)



■ Fairness

- Are threads competing for a lock have a fair chance of acquiring it?

■ Overhead

February 6, 2019

TCCS422: Operating Systems [Winter 2019]
School of Engineering and Technology, University of Washington - Tacoma

L8.34

BUILDING LOCKS

- Locks require hardware support
 - To minimize overhead, ensure fairness and correctness
 - Special “atomic-as a unit” instructions to support lock implementation
 - Atomic-as a unit exchange instruction
 - XCHG
 - Compare and exchange instruction
 - CMPXCHG
 - CMPXCHG8B
 - CMPXCHG16B

February 6, 2019

TCSS422: Operating Systems [Winter 2019]
School of Engineering and Technology, University of Washington - Tacoma

L8.35

HISTORICAL IMPLEMENTATION

- To implement mutual exclusion
 - Disable interrupts upon entering critical sections

```
1  void lock() {  
2      DisableInterrupts();  
3  }  
4  void unlock() {  
5      EnableInterrupts();  
6  }
```

- Any thread could disable system-wide interrupt
 - What if lock is never released?
- On a multiprocessor processor each CPU has its own interrupts
 - Do we disable interrupts for all cores simultaneously?
- While interrupts are disabled, they could be lost
 - If not queued...


February 6, 2019

TCSS422: Operating Systems [Winter 2019]
School of Engineering and Technology, University of Washington - Tacoma

L8.36

SPIN LOCK IMPLEMENTATION

- Operate without atomic-as a *unit* assembly instructions
- “Do-it-yourself” Locks
- Is this lock implementation: (1)Correct? (2)Fair? (3)Performant?



```
1  typedef struct __lock_t { int flag; } lock_t;
2
3  void init(lock_t *mutex) {
4      // 0 → lock is available, 1 → held
5      mutex->flag = 0;
6  }
7
8  void lock(lock_t *mutex) {
9      while (mutex->flag == 1) // TEST the flag
10         ; // spin-wait (do nothing)
11     mutex->flag = 1; // now SET it !
12 }
13
14 void unlock(lock_t *mutex) {
15     mutex->flag = 0;
16 }
```

February 6, 2019

TCCS422: Operating Systems [Winter 2019]
School of Engineering and Technology, University of Washington - Tacoma

L8.37

DIY: CORRECT?

- Correctness requires luck... (e.g. *DIY lock is incorrect*)

Thread1	Thread2
call lock() while (flag == 1) interrupt: switch to Thread 2	
	call lock() while (flag == 1) flag = 1; interrupt: switch to Thread 1
flag = 1; // set flag to 1 (too!)	

- Here both threads have “acquired” the lock simultaneously

February 6, 2019

TCCS422: Operating Systems [Winter 2019]
School of Engineering and Technology, University of Washington - Tacoma

L8.38

DIY: PERFORMANT?

```
void lock(lock_t *mutex)
{
    while (mutex->flag == 1);    // while lock is unavailable, wait...
    mutex->flag = 1;
}
```

- What is wrong with while(<cond>); ?
- Spin-waiting wastes time actively waiting for another thread
- while (1); will “peg” a CPU core at 100%
 - Continuously loops, and evaluates mutex->flag value...
 - Generates heat...

February 6, 2019

TCS422: Operating Systems [Winter 2019]
School of Engineering and Technology, University of Washington - Tacoma

L8.39

TEST-AND-SET INSTRUCTION

- Hardware support required for working locks
- Book presents pseudo code of C implementation
 - TEST-and-SET adds a simple check to the basic spin lock
 - Assumption is this “C code” runs atomically on CPU:

```
1  int TestAndSet(int *ptr, int new) {
2      int old = *ptr;    // fetch old value at ptr
3      *ptr = new;        // store 'new' into ptr
4      return old;        // return the old value
5  }
```

- lock() method checks that TestAndSet doesn't return 1
- Comparison is in the caller
- Can implement the C version (non-atomic) and have some success on a single-core VM

February 6, 2019

TCS422: Operating Systems [Winter 2019]
School of Engineering and Technology, University of Washington - Tacoma

L8.40

DIY: TEST-AND-SET - 2

- C version: requires preemptive scheduler on single core system
- Lock is never released without a context switch
- single-core VM: occasionally will deadlock, doesn't miscount

```
1  typedef struct __lock_t {
2      int flag;
3  } lock_t;
4
5  void init(lock_t *lock) {
6      // 0 indicates that lock is available,
7      // 1 that it is held
8      lock->flag = 0;
9  }
10
11 void lock(lock_t *lock) {
12     while (TestAndSet(&lock->flag, 1) == 1)
13         ; // spin-wait
14 }
15
16 void unlock(lock_t *lock) {
17     lock->flag = 0;
18 }
```

February 6, 2019

TCSS422: Operating Systems [Winter 2019]
School of Engineering and Technology, University of Washington - Tacoma

L8.41

SPIN LOCK EVALUATION

- **Correctness:**
 - Spin locks with atomic Test-and-Set:
Critical sections won't be executed simultaneously by (2) threads
- **Fairness:**
 - No fairness guarantee. Once a thread has a lock, nothing forces it to relinquish it...
- **Performance:**
 - Spin locks perform "busy waiting"
 - Spin locks are best for short periods of waiting (< 1 time quantum)
 - Performance is slow when multiple threads share a CPU
 - Especially if "spinning" for long periods

February 6, 2019

TCSS422: Operating Systems [Winter 2019]
School of Engineering and Technology, University of Washington - Tacoma

L8.42

COMPARE AND SWAP

- Checks that the lock variable has the expected value **FIRST**, before changing its value
 - If so, make assignment
 - Return value at location
- Adds a comparison to TestAndSet
 - Textbook presents C pseudo code
 - Assumption is that the compare-and-swap method runs atomically
- Useful for wait-free synchronization
 - Supports implementation of shared data structures which can be updated atomically (*as a unit*) using the HW support CompareAndSwap instruction
 - Shared data structure updates become “wait-free”
 - Upcoming in Chapter 32

February 6, 2019

TCSS422: Operating Systems [Winter 2019]
School of Engineering and Technology, University of Washington - Tacoma

L8.43

COMPARE AND SWAP

- Compare and Swap

```
1  int CompareAndSwap(int *ptr, int expected, int new) {  
2      int actual = *ptr;  
3      if (actual == expected)  
4          *ptr = new;  
5      return actual;  
6  }
```

- Spin lock

```
1  
2  
3      ; // spin  
4  }
```

**C implementation 1-core VM:
Count is correct, no deadlock**

- X86 provides “**cmpxchgl**” compare-and-exchange instruction
 - **cmpxchg8b**
 - **cmpxchg16b**

February 6, 2019

TCSS422: Operating Systems [Winter 2019]
School of Engineering and Technology, University of Washington - Tacoma

L8.44

TWO MORE “LOCK BUILDING” CPU INSTRUCTIONS

- Cooperative instructions used together to support synchronization on RISC systems
- No support on x86 processors
 - Supported by RISC: Alpha, PowerPC, ARM
- Load-linked (LL)
 - Loads value into register
 - Same as typical load
 - Used as a mechanism to track competition
- Store-conditional (SC)
 - Performs “mutually exclusive” store
 - Allows only one thread to store value

February 6, 2019

TCSS422: Operating Systems [Winter 2019]
School of Engineering and Technology, University of Washington - Tacoma

L8.45

LL/SC LOCK

```
1  int LoadLinked(int *ptr) {
2      return *ptr;
3  }
4
5  int StoreConditional(int *ptr, int value) {
6      if (no one has updated *ptr since the LoadLinked to this address) {
7          *ptr = value;
8          return 1; // success!
9      } else {
10         return 0; // failed to update
11     }
12 }
```

- LL instruction loads pointer value (ptr)
- SC only stores if the load link pointer has not changed
- Requires HW support
 - C code is psuedo code

February 6, 2019

TCSS422: Operating Systems [Winter 2019]
School of Engineering and Technology, University of Washington - Tacoma

L8.46

LL/SC LOCK - 2

```
1 void lock(lock_t *lock) {
2     while (1) {
3         while (LoadLinked(&lock->flag) == 1)
4             ; // spin until it's zero
5         if (StoreConditional(&lock->flag, 1) == 1)
6             return; // if set-it-to-1 was a success: all done
7                     // otherwise: try it all over again
8     }
9 }
10
11 void unlock(lock_t *lock) {
12     lock->flag = 0;
13 }
```

■ Two instruction lock

February 6, 2019

TCSS422: Operating Systems [Winter 2019]
School of Engineering and Technology, University of Washington - Tacoma

L8.47

CHAPTER 29 – LOCK BASED DATA STRUCTURES



February 6, 2019

TCSS422: Operating Systems [Winter 2019]
School of Engineering and Technology, University of Washington - Tacoma

L8.48

OBJECTIVES

- **Chapter 29**
 - **Concurrent Data Structures**
 - **Performance**
 - **Lock Granularity**

February 6, 2019

TCSS422: Operating Systems [Winter 2019]
School of Engineering and Technology, University of Washington - Tacoma

L8.49

LOCK-BASED CONCURRENT DATA STRUCTURES

- **Adding locks to data structures make them thread safe.**
- **Considerations:**
 - **Correctness**
 - **Performance**
 - **Lock granularity**

February 6, 2019

TCSS422: Operating Systems [Winter 2019]
School of Engineering and Technology, University of Washington - Tacoma

L8.50

COUNTER STRUCTURE W/O LOCK

■ Synchronization weary --- not thread safe

```
1  typedef struct __counter_t {
2      int value;
3  } counter_t;
4
5  void init(counter_t *c) {
6      c->value = 0;
7  }
8
9  void increment(counter_t *c) {
10     c->value++;
11 }
12
13 void decrement(counter_t *c) {
14     c->value--;
15 }
16
17 int get(counter_t *c) {
18     return c->value;
19 }
```

February 6, 2019

TCSS422: Operating Systems [Winter 2019]
School of Engineering and Technology, University of Washington - Tacoma

L8.51

CONCURRENT COUNTER

```
1  typedef struct __counter_t {
2      int value;
3      pthread_lock_t lock;
4  } counter_t;
5
6  void init(counter_t *c) {
7      c->value = 0;
8      Pthread_mutex_init(&c->lock, NULL);
9  }
10
11 void increment(counter_t *c) {
12     Pthread_mutex_lock(&c->lock);
13     c->value++;
14     Pthread_mutex_unlock(&c->lock);
15 }
16
```

- Add lock to the counter
- Require lock to change data

February 6, 2019

TCSS422: Operating Systems [Winter 2019]
School of Engineering and Technology, University of Washington - Tacoma

L8.52

CONCURRENT COUNTER - 2

- Decrease counter
- Get value

```
(Cont.)
17  void decrement(counter_t *c) {
18      pthread_mutex_lock(&c->lock);
19      c->value--;
20      pthread_mutex_unlock(&c->lock);
21  }
22
23  int get(counter_t *c) {
24      pthread_mutex_lock(&c->lock);
25      int rc = c->value;
26      pthread_mutex_unlock(&c->lock);
27      return rc;
28  }
```


February 6, 2019

TCSS422: Operating Systems [Winter 2019]
School of Engineering and Technology, University of Washington - Tacoma

L8.53

CONCURRENT COUNTERS - PERFORMANCE

- iMac: four core Intel 2.7 GHz i5 CPU
- Each thread increments counter 1,000,000 times



Traditional vs. sloppy counter
Sloppy Threshold (S) = 1024

Synchronized counter scales poorly.

February 6, 2019

TCSS422: Operating Systems [Winter 2019]
School of Engineering and Technology, University of Washington - Tacoma

L8.54

PERFECT SCALING

- Achieve (N) performance gain with (N) additional resources
- Throughput:
 - Transactions per second
- 1 core
 - N = 100 tps
- 10 core
 - N = 1000 tps

February 6, 2019

TCSS422: Operating Systems [Winter 2019]
School of Engineering and Technology, University of Washington - Tacoma

L8.55

SLOPPY COUNTER

- Provides single logical shared counter
 - Implemented using local counters for each ~CPU core
 - 4 CPU cores = 4 local counters & 1 global counter
 - Local counters are synchronized via local locks
 - Global counter is updated periodically
 - Global counter has lock to protect global counter value
 - Sloppiness threshold (S):
 - Update threshold of global counter with local values
 - Small (S): more updates, more overhead
 - Large (S): fewer updates, more performant, less synchronized
- Why this implementation?
Why do we want counters local to each CPU Core?

February 6, 2019

TCSS422: Operating Systems [Winter 2019]
School of Engineering and Technology, University of Washington - Tacoma

L8.56

